

Exact Schedulability Test for global-EDF Scheduling of Periodic Hard Real-Time Tasks on Identical Multiprocessors

Joël Goossens*

Patrick Meumeu Yomsis†

1 Introduction

Over the years, the preemptive periodic constrained-deadline task model [9] has proven remarkably useful for the modeling of recurring processes that occur in hard real-time computer application systems, where the failure to satisfy any constraint may have disastrous consequences. The problem of scheduling such tasks upon a single processor (CPU) so that all the deadlines are met has been widely studied in the literature and is now well understood. The most important point in this direction being that an *optimal online scheduler*, commonly known as *Earliest Deadline First* (EDF), has been derived. EDF is a priority-based scheduler which assigns priorities to jobs so that the shorter the absolute deadline of a job the higher its priority. This scheduler is optimal with the interpretation that if a periodic constrained-deadline task system can be successfully scheduled with another scheduler upon a single CPU, then it can also be successfully scheduled using EDF. However, a very large number of applications nowadays turns out to be executed upon more than one CPU for practical and economic reasons due to the advent of multicore technologies. For such applications, even though EDF is no longer optimal [8], much recent work gave rise to multiple investigations and thus many alternative algorithms based on this scheduling policy have been developed due to its optimality upon uniprocessor platforms [13]. Most results have been derived under either *global* or *partitioned* scheduling techniques. In *global* scheduling [4], all the tasks are stored in a single priority-ordered queue and the global scheduler selects for execution the highest priority tasks from this queue. In this framework, tasks are allowed to migrate at runtime from one CPU to another in order to complete their executions [6, 3]. In *partitioned* scheduling [2], all the tasks are first assigned statically to the CPUs, then each CPU uses independently its local scheduler at runtime. Despite these two scheduling techniques are *incomparable* [1] in the sense that there are systems which are schedulable with partitioning and not by global and conversely, and despite the high number of interesting results that have already been derived up to now, many open questions still remain to be answered, especially when *global schedulers* are considered. Regarding this kind of schedulers, an important issue consists in deriving an *exact* schedulability test by exploiting on the one

*Brussels University, U.L.B., Brussels, Belgium.

†Postdoctoral researcher of the F.N.R.S., Belgium.

hand the predictability property of the scheduler and by providing on the other hand a feasibility interval so that if it is possible to find a valid schedule for all jobs contained in this interval, then the whole system will be stamped feasible.

Related work. In recent years, as most global schedulers are predictable, extensive efforts have been performed towards addressing the problem of determining a feasibility interval for the global scheduling of periodic constrained-deadline tasks upon multiprocessor platforms. That is, to derive an interval of time so that if it is possible to find a valid schedule for all jobs contained in this interval, then the whole system is feasible. Up to now, sound results have been obtained only in the particular case where tasks are scheduled by using an Fixed-Task-Priority (FTP) scheduler [6, 7]. Being an FTP scheduler one where all the jobs belonging to a task are assigned the same priority as the priority assigned to the task beforehand (i.e., at design time). We are not currently aware of any existing result concerning the feasibility interval for Fixed-Job-Priority (FJP) schedulers in the literature, except the one proposed by Leung in [12]. However, we show that this result is actually wrong. An FJP scheduler is one where two jobs belonging to the same task may be assigned different priorities.

This research. In this paper, we derive a feasibility interval for an FJP scheduler, namely global-EDF. To the best of our knowledge, this will be the *first* valuable feasibility interval for FJP schedulers since the one proposed by Leung in [12] is flawed. Based on this feasibility interval and considering the predictability property of this scheduler, our main contribution is therefore an *exact* schedulability test for the global-EDF scheduling of periodic hard real-time tasks upon identical multiprocessor platforms.

Paper organization. The remainder of this paper is structured as follows. Section 2 presents the system model and the scheduler that are used throughout the paper. Section 3 provides the reader with some useful definitions and properties. Section 4 presents our main contribution. Finally, Section 5 concludes the paper.

2 System model

Throughout this paper, all timing characteristics in our model are assumed to be non-negative integers, i.e., they are multiples of some elementary time interval (for example the CPU tick, the smallest indivisible CPU time unit).

2.1 Task specifications

We consider the preemptive scheduling of a hard real-time system $\tau \stackrel{\text{def}}{=} \{\tau_1, \tau_2, \dots, \tau_n\}$ composed of n tasks upon m identical CPUs according to the following interpretations.

- ▷ *Preemptive scheduling:* an executing task may be interrupted at any instant in time and have its execution resumed later.

▷ *Identical CPUs*: all the CPUs have the same computing capacities.

Each task τ_i is a periodic constrained-deadline task characterized by four parameters (O_i, C_i, D_i, T_i) where O_i is the first release time (offset), C_i is the Worst Case Execution Time (WCET), $D_i \leq T_i$ is the relative deadline and T_i is the period, i.e., the *exact* inter-arrival time between two consecutive releases of task τ_i . These parameters are given with the interpretation that task τ_i generates an infinite number of successive jobs $\tau_{i,j}$ from time instant O_i , with execution requirement of at most C_i each, the j^{th} job which is released at time $O_{i,j} \stackrel{\text{def}}{=} O_i + (j - 1) \cdot T_i$ must complete within $[O_{i,j}, d_{i,j})$ where $d_{i,j} \stackrel{\text{def}}{=} O_{i,j} + D_i$, the absolute deadline of job $\tau_{i,j}$.

We assume without any loss of generality that $O_i \geq 0$, $\forall i \in \{1, 2, \dots, n\}$ and we denote by O_{\max} the maximal value among all task offsets, i.e., $O_{\max} \stackrel{\text{def}}{=} \max\{O_1, O_2, \dots, O_n\}$. We denote by P the *hyperperiod* of the system, i.e., the least common multiple (lcm) of all tasks periods: $P \stackrel{\text{def}}{=} \text{lcm}\{T_1, T_2, \dots, T_n\}$. Also, we denote by C_τ the sum of the WCETs of all tasks in τ : $C_\tau \stackrel{\text{def}}{=} \sum_{i=1}^n C_i$.

Job $\tau_{i,j}$ is said to be *active* at time t if and only if $O_{i,j} \leq t$ and $\tau_{i,j}$ is not completed yet. More precisely, an active job is said to be *running* at time t if it has been allocated to a CPU and is being executed. Otherwise, the active job is said to be *ready* and is in the ready queue of the operating system.

We assume that all the tasks are independent, i.e., there is no communication, no precedence constraint and no shared resource (except for the CPUs) between tasks. Also, we assume that any job $\tau_{i,j}$ cannot be executed in parallel, i.e., no job can execute upon more than one CPU at any instant in time.

2.2 Scheduler specifications

We consider that tasks are scheduled by using the Fixed-Job-Priority (FJP) scheduler global-EDF. That is, the following two properties are always satisfied: (i) the shorter the absolute deadline of a job the higher its priority and (ii) a job may begin execution on any CPU and a preempted job may resume execution on the same CPU as, or a different CPU from, the one it had been executing on prior to preemption. We assume in this research that the preemptions and migrations of all tasks and jobs in the system are allowed at no cost or penalty.

3 Definitions and properties

In this section we provide definitions and properties that will help us establishing our exact schedulability test. First, we formalize the notions of *synchronous* and *asynchronous* systems, *schedule* and *valid schedule*, and *configuration*.

Definition 1 ((A)Synchronous systems). *A task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$ is said to be synchronous if each task in τ has its first job released at the same time-instant c , i.e., $O_i = c$ for all $1 \leq i \leq n$. Otherwise, τ is said to be asynchronous.*

Definition 2 (Schedule $\sigma(t)$). For any task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$ and any set of m identical CPUs $\{\pi_1, \pi_2, \dots, \pi_m\}$, the schedule $\sigma(t)$ of system τ at time-instant t is defined as $\sigma : \mathbb{N} \rightarrow \{1, 2, \dots, n\}^m$ where $\sigma(t) \stackrel{\text{def}}{=} (\sigma_1(t), \sigma_2(t), \dots, \sigma_m(t))$ with

$$\sigma_j(t) \stackrel{\text{def}}{=} \begin{cases} 0, & \text{if there is no task scheduled on } \pi_j \text{ at time-instant } t \\ i, & \text{if task } \tau_i \text{ is scheduled on } \pi_j \text{ at time-instant } t. \end{cases}$$

Definition 3 (Valid schedule). A schedule σ of a task system $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$ is said to be valid if and only if no task in τ ever misses a deadline when tasks are released at their specified released times.

Definition 4 (Configuration $C_S(\tau, t)$). Let S be the schedule of a task system τ . We define the configuration of the schedule S at time t , denoted by $C_S(\tau, t)$, to be the n -tuple $(e_{1,t}, e_{2,t}, \dots, e_{n,t})$, where $e_{i,t}$ is the amount of time for which task τ_i has executed since its last release time up until time t , and $e_{i,t}$ is undefined if $t < O_i$. In the latter case, $C_S(\tau, t)$ is undefined.

Following Definition 4, the configuration $C_S(\tau, t)$ at time t of a schedule S is defined if and only if $t \geq O_{\max}$. Moreover, we have $0 \leq e_{i,t} \leq C_i, \forall t \geq 0$. Now let $t \geq O_{\max}$ and $t' \geq O_{\max}$ be two time instants such that $t \neq t'$, we denote by $C_S(\tau, t) \succeq C_S(\tau, t')$ the fact that $e_{i,t} \geq e_{i,t'}, \forall 1 \leq i \leq n$.

From now on, we always assume an implementation of global-EDF which is *deterministic, work-conserving* and *request-dependent* [6] according to the following definitions.

Definition 5 (Deterministic schedulers). A scheduler is said to be deterministic if and only if it generates a unique schedule for any given set of jobs.

Definition 6 (Work-conserving schedulers). A scheduler is said to be work-conserving if and only if it never idles a CPU while there is at least one active ready task.

Definition 7 (Request-dependent schedulers). A scheduler is said to be request-dependent if and only if for any two tasks $\tau_i, \tau_j \in \tau$ and any two jobs $\tau_{i,k}, \tau_{j,\ell}$ such that $\tau_{i,k}$ is assigned a higher priority than $\tau_{j,\ell}$, then we also have that $\tau_{i,k+P/T_i}$ is assigned a higher priority than $\tau_{j,\ell+P/T_j}$.

Informally speaking, Definition 7 requires that the very same total order is used each hyperperiod between the “corresponding” jobs in terms of priorities.

The deterministic and request-dependent requirements are mandatory to ensure a periodic schedule and these requirements impact, eventually, on the global-EDF tie-breaker, i.e., the tie-breaker must be deterministic and request-dependent.

Some further definitions.

Definition 8 (\mathcal{A} -feasibility). A periodic constrained-deadline task system τ is said to be \mathcal{A} -schedulable upon a set of m identical CPUs if all the tasks in τ meet all their deadlines when scheduled using scheduler \mathcal{A} , i.e., scheduler \mathcal{A} produces a valid schedule.

Definition 9 (Predictability). *A scheduler \mathcal{A} is said to be predictable if the \mathcal{A} -feasibility of a set of tasks implies the \mathcal{A} -feasibility of another set of tasks with identical release times and deadlines, but smaller execution requirements.*

Before we present the main result of this paper, we need to introduce the following notations and results taken from [10] and [6].

Lemma 1 (Ha and Liu [10]). *Any work-conserving and FJP scheduler is predictable upon identical multiprocessor platforms.*

Thanks to Lemma 1, we are guaranteed that the global-EDF scheduler is predictable. Indeed, global-EDF is a work-conserving and FJP scheduler. Thereby, given a periodic constrained-deadline task system τ , we can always assume an instance of τ in which all jobs execute for their whole WCETs. This leads us to consider hereafter a system having known jobs release times, deadlines and execution times. If $\mathcal{S}^{\text{worst}}$ is the valid schedule obtained with these parameters by using the global-EDF scheduler, then we are guaranteed to successfully schedule every other possible instance of τ in which jobs can execute for less than their WCETs by using the same global-EDF scheduler.

Lemma 2 (Cucu and Goossens [6]). *Let \mathcal{S} be the schedule of a periodic constrained-deadline task system τ constructed by using the global-EDF scheduler. If the deadlines of all task computations are met, then \mathcal{S} is periodic from some point with a period equal to P .*

Lemma 3 (Inspired from Cucu and Goossens [7]). *Let \mathcal{S} be the schedule of a periodic constrained-deadline task system τ constructed by using the global-EDF scheduler. Then, for each task τ_i and for each time instant $t_1 \geq O_i$, we have $e_{i,t_1} \geq e_{i,t_2}$, where $t_2 \stackrel{\text{def}}{=} t_1 + P$.*

Proof. The proof is made by contradiction. We assume there is some task τ_{j_1} and some time instant $t_1 \geq O_{j_1}$ such that $e_{j_1,t_1} < e_{j_1,t_2}$, where $t_2 = t_1 + P$. Then there must be some time instant $t'_1 < t_1$ such that τ_{j_1} is active at both t'_1 and $t'_2 = t'_1 + P$, and τ_{j_1} is scheduled at t'_2 while is not at t'_1 . This can only occur if there is another task τ_{j_2} , which is active (running) at t'_1 but not at t'_2 . But this implies that $e_{j_2,t'_1} < e_{j_2,t'_2}$. Thus, we may repeat the above argument to produce an infinite progression of tasks $\tau_{j_3}, \tau_{j_4}, \dots$, for which no lower bound will exist for the time at which the tasks in the sequence are active. But this is impossible, since every task τ_i in τ has an initial release time O_i and the time is discrete in our model of computation. The lemma follows. \square

Note that for any task $\tau_i \in \tau$ and for any time instant $t \geq O_{\max}$, $C_{\mathcal{S}}(\tau, t)$ is *monotonically decreasing* relative to t with period P thanks to Lemma 3, i.e., $C_{\mathcal{S}}(\tau, t + k \cdot P) \succeq C_{\mathcal{S}}(\tau, t + (k + 1) \cdot P) \quad \forall k \in \mathbb{N}$

4 Exact schedulability test

In this section we provide an exact schedulability test for the global-EDF scheduling of periodic hard real-time tasks upon identical multiprocessor platforms. It is worth noticing that we assume in this section that each job of the same

task (say τ_i) has an execution requirement which is exactly C_i time units thanks to the predictability property of this scheduler. Based on the later result, the intuitive idea behind our approach is to construct a schedule by using an implementation of global-EDF which follows hypothesis described in Section 3, then check to see if the deadlines of all task computations are met. However, for this method to work we need to establish an “a priori” time interval within which we need to construct the schedule. If the task system τ is synchronous, then such a time interval is known: $[0, P)$ where $P = \text{lcm}\{T_1, T_2, \dots, T_n\}$ see [6] for details. Unfortunately, if the task system τ is asynchronous, such a time interval is unknown, in the following we will fill the gap.

As the task system τ is composed of periodic tasks, the idea thereby consists in simulating the system until the schedule becomes periodic, i.e., the *steady phase* representing the general timely behavior of the system from a certain time instant is reached. This steady phase is reached when two configurations separated by P time units are identical.

Before going any further in this paper, it is worth noticing the following interesting observations.

Observation 1. By extending the results obtained in the uniprocessor framework to the multiprocessor platforms, Leung claimed in [12] that an exact feasibility condition for global-EDF consists in checking if (i) every deadline is met until time $O_{\max} + 2P$ and (ii) the configurations at instants $O_{\max} + P$ and $O_{\max} + 2P$ are identical. Anyway, this is a flaw, since there are schedulable task systems that reach their steady phase *later* than $O_{\max} + 2P$, as shown by Counterexample 1 taken from [5].

Counterexample 1 (Braun and Cucu [5]). *Consider the following periodic task system: $\tau_1 = (O_1 = 0, C_1 = 2, D_1 = T_1 = 3)$, $\tau_2 = (O_2 = 4, C_2 = 3, D_2 = T_2 = 4)$, $\tau_3 = (O_3 = 1, C_3 = 3, D_3 = T_3 = 6)$ to be scheduled with global-EDF upon $m = 2$ CPUs.*

By building the schedule (see Figure 1), it is possible to see that at time $O_{\max} + P = 4 + 12 = 16$ and $O_{\max} + 2P = 4 + 2 \cdot 12 = 28$ the steady phase has not yet been reached (at times 17 and 29 there are two different configurations). However, the steady phase is reached after a further hyperperiod. Since no deadline is missed, the task system is schedulable with global-EDF.

The flaws in the results proposed by Leung in [12] come from many sources. The paper was actually centered on the Least Laxity First (LLF) scheduler defined as follows.

Definition 10 (LLF scheduler). *The LLF scheduler always executes the jobs with least laxity; being the laxity of a job its absolute deadline minus the sum of its remaining processing time and the current time.*

▷ **Flaw 1.** The first flaw is about the comparison between the LLF scheduler and the global-EDF scheduler. Indeed it was claimed in [12] that every instance of jobs schedulable by global-EDF upon m CPUs is also schedulable by LLF upon m CPUs. This flaw, presented as one of the main results of the paper, has been pointed out by Kalyanasundaram et al. in [11].

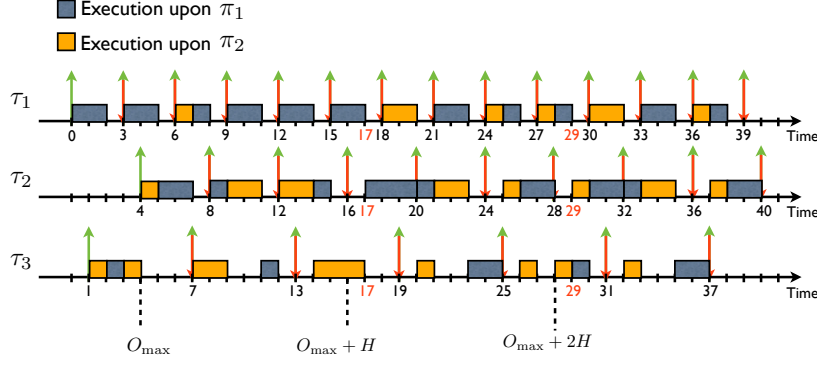


Figure 1: $[0, O_{\max} + 2P]$ is not a feasibility interval for the global-EDF scheduler.

▷ **Flaw 2.** The second flaw lies in the proof of Lemma 2 of Leung’s paper replicated here (except we consider global-EDF).

Lemma 4 (Lemma 2, pages 216–217 of [12]). *Let S be the schedule of a task system τ constructed by using the global-EDF scheduler upon $m \geq 1$ CPUs. If τ is schedulable by using the global-EDF scheduler on m CPUs, then $C_S(\tau, t_1) = C_S(\tau, t_2)$ where $t_1 \stackrel{\text{def}}{=} O_{\max} + P$ and $t_2 \stackrel{\text{def}}{=} t_1 + P$.*

The argument stating that:

the m CPUs are always busy in the interval $[t_1, t_2]$

is incorrect; this is a *uniprocessor* argument not valid in a multiprocessor context. Indeed, considering Counterexample 1, it is not difficult to see in Figure 1 that $t_1 = 16$, $t_2 = 28$ and $C_S(\tau, 16) \neq C_S(\tau, 28)$. However, in the time-slots $[17, 18]$ and $[23, 24]$, only one CPU (here, CPU π_1) out of two is actually busy by the execution of the jobs.

Observation 2. Although we address the global-EDF scheduling problem of periodic constrained-deadline task systems, Counterexample 2 give evidence of possibly late occurrence of the *steady phase* in the valid schedule S of a task system τ . Indeed, it shows that the steady phase can be reached after a time-instant as large as $O_{\max} + 42 \cdot P$.

Counterexample 2. Consider the following periodic task system: $\tau_1 = (O_1 = 225, C_1 = 90, D_1 = T_1 = 161)$, $\tau_2 = (O_2 = 115, C_2 = 40, D_2 = T_2 = 161)$, $\tau_3 = (O_3 = 0, C_3 = 72, D_3 = T_3 = 161)$, $\tau_4 = (O_4 = 129, C_4 = 120, D_4 = T_4 = 161)$ to be scheduled with global-EDF upon $m = 2$ CPUs.

By building the schedule using an open source simulation tool such as STORM¹ (we implemented a deterministic and request-dependent EDF tie-breaker), it is possible to see that at time-instants $O_{\max} + 42P = 6987$ and $O_{\max} + 43P = 7148$ the steady phase has not been reached yet (there are two different configurations at time-instants 6988

¹STORM stands for “Simulation Tool for Real-time Multiprocessor Scheduling Evaluation” and is a simulation tool developed at Ircyn, École Centrale de Nantes, France.

and 7149). However, the steady phase is reached after a further hyperperiod. Again, since no deadline is missed, the task system is schedulable with global-EDF.

It thus follows from Lemma 2, Observation 1 and Observation 2 the conjecture that integer $k \in \mathbb{N}^+$ in Expression $(O_{\max} + k \cdot P)$ for the time-instant to reach the steady phase must be a function of tasks parameters.

Lemma 5. *Let \mathcal{S} be the valid schedule of an asynchronous periodic constrained-deadline task system τ . We assume that \mathcal{S} has been constructed by using the global-EDF scheduler. Then an upper bound on the time-instant up to which there exists $t > 0$ such that $C_{\mathcal{S}}(\tau, t - P) = C_{\mathcal{S}}(\tau, t)$ is given by $t_{\text{up}} \stackrel{\text{def}}{=} O_{\max} + (C_{\tau} + 1) \cdot P$*

Proof. Let \mathcal{S} be the valid schedule of an asynchronous periodic constrained-deadline task system τ , constructed by using the global-EDF scheduler. By definition, O_{\max} is the first time-instant at which all tasks in τ are released, the steady state cannot start before that time-instant. Now we will upper-bound the first time-instant where the schedule starts to repeat. The worst-case scenario is the one where there is a *different* system configuration, between two successive hyperperiods, a *maximal* number of times starting from O_{\max} . I.e., $e_{i, O_{\max}} = C_i \quad \forall i$ and $C_{\mathcal{S}}(\tau, O_{\max} + k \cdot P) \neq C_{\mathcal{S}}(\tau, O_{\max} + (k+1) \cdot P) \quad \forall k \in [0, \hat{k}]$ and $C_{\mathcal{S}}(\tau, O_{\max} + \hat{k} \cdot P) = C_{\mathcal{S}}(\tau, O_{\max} + (\hat{k}+1) \cdot P)$. By Lemma 3 and the fact that time is discrete in our model of computation, the worst-case scenario—i.e., the scenario which maximizes \hat{k} —corresponds to the case where each time the system configuration differs from the previous one in the following way: $\exists \ell \in [1, n]$ such that: $e_{i, O_{\max} + (k+1) \cdot P} = e_{i, O_{\max} + k \cdot P} \quad \forall i \neq \ell$ and $e_{\ell, O_{\max} + (k+1) \cdot P} = e_{\ell, O_{\max} + k \cdot P} - 1$. Since $0 \leq e_{i, t} \leq C_i$ an upper-bound for \hat{k} is $\sum_{i=1}^n C_i$. And by Lemma 2 at time $t_{\text{up}} = O_{\max} + (\sum_{i=1}^n C_i + 1) \cdot P = O_{\max} + (C_{\tau} + 1) \cdot P$ we have $C_{\mathcal{S}}(\tau, t_{\text{up}} - P) = C_{\mathcal{S}}(\tau, t_{\text{up}})$ and the schedule is periodic with a period of P . \square

Now we have the material to define an “exact” schedulability test for the global-EDF scheduling of periodic hard real-time tasks upon identical multiprocessor platforms.

Exact Schedulability Test. Let τ be an asynchronous periodic constrained-deadline task system. Let π be a platform consisting of m identical CPUs. We assume that τ is scheduled upon π by using the global-EDF scheduler. Then τ is schedulable if and only if (i) all deadlines are met in $[0, t_{\text{up}})$ and (ii) $C_{\mathcal{S}}(\tau, t_{\text{up}} - P) = C_{\mathcal{S}}(\tau, t_{\text{up}})$ where t_{up} is defined in Lemma 5.

Proof. If there is a deadline miss in $[0, t_{\text{up}})$, then the system is clearly not schedulable. Otherwise, this schedulability test is a direct consequence of Lemma 5, Lemma 2 (the periodicity property of the schedule) and Lemma 1 (the predictability property of the global-EDF scheduler). \square

5 Conclusion

In this paper, we considered the scheduling problem of hard real-time systems composed of periodic constrained-deadline tasks upon identical multiprocessor platforms. We assumed that tasks were scheduled by using the global-EDF

scheduler and we provided an *exact* schedulability test for this scheduler. Also, we showed by means of a counterexample that the feasibility interval, and thus the schedulability test, proposed by Leung [12] is incorrect and we showed which arguments are actually incorrect.

References

- [1] S. Baruah. Techniques for multiprocessor global schedulability analysis. In *Real-Time Systems Symposium*, pages 119–128, 2007.
- [2] S. Baruah and N. Fisher. The partitioned multiprocessor scheduling of deadline-constrained sporadic task systems. *IEEE Trans. on Computers* 55 (7), pages 918–923, 2006.
- [3] M. Bertogna. *Real-Time Scheduling Analysis for Multiprocessor Platforms*. PhD thesis, Scuola Superiore Sant’Anna, Pisa, 2008.
- [4] M. Bertogna, M. Cirinei, and G. Lipari. Schedulability analysis of global scheduling algorithms on multiprocessor platforms. *IEEE Trans. Parallel Distrib. Syst.*, 20(4):553–566, 2009.
- [5] C. Braun and L. Cucu. Negative results on idle intervals and periodicity for multiprocessor scheduling under EDF. In *Junior Researcher Workshop on Real-Time Computing*, 2007.
- [6] L. Cucu and J. Goossens. Feasibility intervals for fixed-priority real-time scheduling on uniform multiprocessors. In *the 11th IEEE International Conference on Emerging Technologies and Factory Automation*, pages 397–405, 2006.
- [7] L. Cucu and J. Goossens. Feasibility intervals for multiprocessor fixed-priority scheduling of arbitrary deadline periodic systems. In *Design Automation and Test in Europe*, pages 1635–1640, 2007.
- [8] S. K. Dhall and C. L. Liu. On a real-time scheduling problem. *Operations Research*, 26:127–140, 1978.
- [9] J. Goossens, S. Funk, and S. Baruah. Priority-driven scheduling of periodic task systems on multiprocessors. *Real-time Systems Journal*, 25:187–205, 2003.
- [10] R. Ha and J. Liu. Validating timing constraints in multiprocessor and distributed real-time systems. In *14th IEEE International Conference on Distributed Computing Systems*, pages 162–171, 1994.
- [11] B. Kalyanasundaram, K. R. Pruhs, and E. Torng. Errata: A new algorithm for scheduling periodic, real-time tasks. *Algorithmica*, 28:269–270, 2000.
- [12] J. Y.-T. Leung. A new algorithm for scheduling periodic real-time tasks. *Algorithmica*, 4:209–219, 1989.
- [13] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20(1):46–61, 1973.